

# On the Impact of Identifiers on Local Decision<sup>\*</sup>

Pierre Fraigniaud<sup>1\*\*</sup>, Magnús M. Halldórsson<sup>2\*\*\*</sup>, and Amos Korman<sup>\*\*</sup>

<sup>1</sup> CNRS and University Paris Diderot, France.

<sup>2</sup> ICE-TCS, School of Computer Science, Reykjavik University, Iceland.

**Abstract.** The issue of identifiers is crucial in distributed computing. Informally, identities are used for tackling two of the fundamental difficulties that are inherent to deterministic distributed computing, namely: (1) *symmetry breaking*, and (2) *topological information gathering*. In the context of *local computation*, i.e., when nodes can gather information only from nodes at bounded distances, some insight regarding the role of identities has been established. For instance, it was shown that, for large classes of *construction* problems, the role of the identities can be rather small. However, for the identities to play no role, some other kinds of mechanisms for breaking symmetry must be employed, such as edge-labeling or sense of direction. When it comes to local distributed *decision* problems, the specification of the decision task does not seem to involve symmetry breaking. Therefore, it is expected that, assuming nodes can gather sufficient information about their neighborhood, one could get rid of the identities, without employing extra mechanisms for breaking symmetry. We tackle this question in the framework of the *LOCAL* model. Let LD be the class of all problems that can be *decided* in a constant number of rounds in the *LOCAL* model. Similarly, let LD<sup>\*</sup> be the class of all problems that can be decided at constant cost in the anonymous variant of the *LOCAL* model, in which nodes have no identities, but each node can get access to the (anonymous) ball of radius  $t$  around it, for any  $t$ , at a cost of  $t$ . It is clear that LD<sup>\*</sup>  $\subseteq$  LD. We conjecture that LD<sup>\*</sup> = LD. In this paper, we give several evidences supporting this conjecture. In particular, we show that it holds for *hereditary* problems, as well as when the nodes know an arbitrary upper bound on the total number of nodes. Moreover, we prove that the conjecture holds in the context of *non-deterministic* local decision, where nodes are given certificates (independent of the identities, if they exist), and the decision consists in verifying these certificates. In short, we prove that NLD<sup>\*</sup> = NLD.

**Keywords:** Distributed complexity; locality; identities; decision problems; symmetry breaking; non-determinism.

---

<sup>\*</sup> This work is supported by the *Jules Verne* Franco-Icelandic bilateral scientific framework.

<sup>\*\*</sup> E-mail: {pierre.fraigniaud,amos.korman}@liafa.univ-paris-diderot.fr. Additional support from ANR project DISPLEXITY, and INRIA project GANG.

<sup>\*\*\*</sup> Supported by Iceland Research Foundation grant-of-excellence 90032021. E-mail: mmh@ru.is.

# 1 Introduction

## 1.1 Background and Motivation

The issue of identifiers is crucial in distributed computing [2,33]. Indeed, the correct operation of deterministic protocols often relies on the assumption that each processor  $u$  comes with a unique *identity*,  $\text{Id}(u)$  [9]. Informally, in network computing, such an identity assignment is crucial for tackling two of the fundamental difficulties that are inherent to distributed computing, namely: (1) *symmetry breaking*, and (2) *topological information gathering*.

The use of identities for tackling the above two difficulties is illustrated well in the context of *local* algorithms [30,32]. Indeed, in the  $\mathcal{LOCAL}$  model [35], an algorithm that runs in  $t$  communication rounds, assuming an identity assignment, can be viewed as composed of two parts: first, collecting at each node  $u$ , the ball  $B(u, t)$  of radius  $t$  around it (together with the inputs of nodes), and second, deciding the output at  $u$  based solely on the information in  $B(u, t)$ . To achieve these two tasks, one should first obtain the ball  $B(u, t)$ , which may not be possible if the underlying graph is anonymous (i.e., without identities). Moreover, even if obtaining the ball is possible, e.g., if the structure of the graph allows it, the absence of unique identities given to the nodes may prevent the algorithm from breaking symmetry. For example, in the absence of unique identities, it is impossible to design a distributed deterministic coloring algorithm, even for the symmetric connected graph composed of two nodes only. In fact, to the best of our knowledge, all algorithms in the  $\mathcal{LOCAL}$  model are designed assuming the presence of pairwise distinct identities or some other type of node-labeling or edge-labeling, including, e.g., sense of direction [5,22,29,31,33,34].

The seminal paper of Naor and Stockmeyer [33] provides an important insight regarding the role of identities in local computation. Informally, they show that, even though identities are necessary, in many cases the actual values of identities is not crucial, and only their relative order matters. Specifically, [33] shows that for a particular class of problems, called LCL (for *Locally Checkable Languages*), if there exists a local algorithm that, for any identity assignment, constructs an instance of a problem in LCL in constant number of rounds, then there exists an *order invariant*<sup>3</sup> algorithm for that problem that runs in the same number of rounds. LCL restricts its concern to graphs with constant maximum degree, and to problems with a constant number of inputs. The assumption on the size of the inputs of problems in LCL was shown necessary in [21], by exhibiting a natural problem that is locally checkable, has unbounded input size, can be solved in 1 round with identities, but cannot be solved in constant time by any order invariant algorithm. The role of identities can also be gauged by comparing their impact to that of “orientation mechanisms”. For instance, Göös et al. [20] have shown that for a large class of optimization problems, called PO-checkable

---

<sup>3</sup> Essentially, an order invariant algorithm uses the actual values of the identities only to impose an ordering between the nodes, that is, it behaves the same for any two identity assignments that preserve the total order between the nodes. For more details refer to [33].

problems, local algorithms do not benefit from any kind of identifiers: if a PO-checkable optimization problem can be approximated with a local algorithm, the same approximation factor can be achieved in anonymous networks if the network is provided with a port-numbering and an orientation.

The discussion above involved distributed *construction* tasks, including, e.g., graph coloring [5,29,30,33,34], maximal independent set [30,34], and maximal matching [22,31]. When it comes to distributed *decision* tasks [13,14], symmetry breaking issues do not however seem to play a role. Informally, a decision task requires the nodes to “collectively decide” whether the given instance (i.e., a graph with inputs to the nodes) satisfies some specific properties. For instance, deciding coloring requires, given a colored graph, to check whether this graph is properly colored. The meaning of “collectively decide” is as follows. On a legal instance, all nodes should output “yes”, and on an illegal one, at least one node should output “no”. Note that it is not really important whether this node is unique or not; hence, this specification does not inherently require any symmetry breaking. Therefore, assuming that each node  $u$  can obtain the ball  $B(u, t)$ , it makes sense that the assumption of having an identity assignment may not be crucial for achieving correct decision.

## 1.2 Model and Objectives

We tackle the question of whether identities play a role in decision problems in the framework of the aforementioned *LOCAL* model [35], which is a standard distributed computing model capturing the essence of locality. Recall that, in this model, processors are nodes of a connected network  $G = (V(G), E(G))$ , have pairwise distinct identities, and have inputs. More formally, a *configuration* is a triplet  $(G, \mathbf{x}, \text{Id})$  where  $G$  is a connected graph, every node  $v \in V(G)$  is assigned as its *local input* a binary string  $\mathbf{x}(v) \in \{0, 1\}^*$ , and  $\text{Id}(v)$  denotes the identity of node  $v$ . (In some problems, the local input of every node is empty, i.e.,  $\mathbf{x}(v) = \epsilon$  for every  $v \in V(G)$ , where  $\epsilon$  denotes the empty binary string). Processors are woken up simultaneously, and computation proceeds over the input configuration  $(G, \mathbf{x}, \text{Id})$  in fault-free synchronous *rounds* during which every processor exchanges messages of unlimited size with its neighbors in the underlying network  $G$ , and performs arbitrary individual computations on its data. In many cases, the running time of an algorithm is measured with respect to the size  $n$  of  $G$ : the running time of an algorithm is defined as the maximum number of rounds it takes to terminate at all nodes, over all possible  $n$ -node networks. Similarly to [21,33], we consider algorithms whose running time is independent of the size of the network, that is they run in constant time.

Let  $B(u, t)$  be the ball centered at  $u$ , of radius  $t$ , excluding the edges between two nodes at distance exactly  $t$  from  $u$ . As mentioned before, without loss of generality, any algorithm running in time  $t = O(1)$  in the *LOCAL* model consists of:

1. Collecting (in  $t$  rounds) at every node  $u$  the structure of the ball  $B(u, t)$  together with all the inputs  $\mathbf{x}(v)$  and identities  $\text{Id}(v)$  of these nodes, and,

2. Performing some individual computation at every node.

We define the *anonymous LOCAL* model similarly to the *LOCAL* model, except that nodes have no identities. More precisely, an input configuration in the anonymous *LOCAL* model is just a pair  $(G, \mathbf{x})$ . An algorithm running in time  $t = O(1)$  in the anonymous *LOCAL* model consists of:

1. Getting at every node  $u$  a snapshot of the structure of the ball  $B(u, t)$  together with all the inputs of the nodes in this ball, and,
2. Performing some individual computation at every node.

Note that the *anonymous LOCAL* model does not explicitly involve communications between nodes. Instead, it implicitly assumes that the underlying network supports the snapshot operation. Clearly, this model is not stronger than the *LOCAL* model, and possibly even strictly weaker, since a node  $u$  can no longer base its individual computation on the identities of the nodes in the ball  $B(u, t)$ . One can think of various other “anonymous” models, i.e., which do not involve node identities. In particular, there is a large literature on distributed computing in networks without node identities, where symmetry breaking is enabled thanks to locally disjoint port numbers (see, e.g., [18]). We consider the anonymous *LOCAL* model to isolate the role of node identities from other symmetry breaking mechanisms.<sup>4</sup> Our aim is to compare the power of the anonymous *LOCAL* model with the standard *LOCAL* model in order to capture the impact of identities on local distributed decision.

Recall from [13] that a *distributed language* is a decidable collection  $\mathcal{L}$  of configurations. (Since an undecidable collection of configurations remains undecidable in the distributed setting too, we consider only decidable collections of configurations). A typical example of a language is

$$\text{Coloring} = \{(G, \mathbf{x}) \mid \forall v \in V(G), \forall w \in N(v), \mathbf{x}(v) \neq \mathbf{x}(w)\} ,$$

where  $N(v)$  denotes the (open) neighborhood of  $v$ , that is, all nodes at distance exactly 1 from  $v$ . Still following the terminology from [13], we say that a distributed algorithm  $A$  *decides* a distributed language  $\mathcal{L}$  if and only if for every configuration  $(G, \mathbf{x})$ , every node of  $G$  eventually terminates and outputs “yes” or “no”, satisfying the following decision rules:

- if  $(G, \mathbf{x}) \in \mathcal{L}$ , then each node outputs “yes”;
- if  $(G, \mathbf{x}) \notin \mathcal{L}$ , then at least one node outputs “no”.

In the (non-anonymous) *LOCAL* model, these two rules must be satisfied for every identity assignment. That is, all processes must output “yes” on a legal instance, independent of their identities. And, on an illegal instance, at least

---

<sup>4</sup> In some sense, the anonymous *LOCAL* model is the strongest model among all models without node identities. Indeed, there are network problems that can be solved in the anonymous *LOCAL* model which cannot be solved in the aforementioned model that is based on locally disjoint port numbers. A simple example is to locally detect the absence of a 3-node cycle.

one node must output “no”, for every identity assignment. Note that this node may potentially differ according to the identity assignment. Some languages can be decided in constant time (e.g., **Coloring**), while others can easily be shown not to be decidable in constant time (e.g., “is the network planar?”). In contrast to the above examples, there are some languages whose status is unclear. To elaborate on this, consider the particular case where it is required to decide whether the network belongs to some specified family  $\mathcal{F}$  of graphs. If this question can be decided in a constant number of communication rounds, then this means, informally, that the family  $\mathcal{F}$  can somehow be characterized by relatively simple conditions. For example, a family  $\mathcal{F}$  of graphs that can be characterized as consisting of all graphs having no subgraph from  $\mathcal{C}$ , where  $\mathcal{C}$  is some specified finite set of graphs, is obviously decidable in constant time. However, the question of whether a family of graphs can be characterized as above is often non-trivial. For example, characterizing cographs as precisely the graphs with no induced  $P_4$ , attributed to Seinsche [36], is not easy, and requires nontrivial usage of modular decomposition.

We are now ready to define one of our main subjects of interest, the classes LD and LD\*. Specifically, LD (for *local decision*) is the class of all distributed languages that can be decided by a distributed algorithm that runs in a constant number of rounds in the  $\mathcal{LOCAL}$  model [13]. Similarly, LD\*, the anonymous version of LD, is the class of all distributed languages that can be decided by a distributed algorithm that runs in a constant number of rounds in the anonymous  $\mathcal{LOCAL}$  model. By definition,  $\text{LD}^* \subseteq \text{LD}$ . We conjecture that

$$\text{LD}^* = \text{LD}.$$

In this paper, we provide several evidences supporting this conjecture. In addition, we investigate the *non-deterministic* version of these classes, and prove that they coincide. More specifically, a distributed *verification* algorithm is a distributed algorithm  $A$  that gets as input, in addition to a configuration  $(G, \mathbf{x})$ , a global *certificate vector*  $\mathbf{y}$ , i.e., every node  $v$  of a graph  $G$  gets as input two binary strings, an input  $\mathbf{x}(v) \in \{0, 1\}^*$  and a certificate  $\mathbf{y}(v) \in \{0, 1\}^*$ . A verification algorithm  $A$  verifies  $\mathcal{L}$  if and only if for every input configuration  $(G, \mathbf{x})$ , the following hold:

- if  $(G, \mathbf{x}) \in \mathcal{L}$ , then there exists a certificate  $\mathbf{y}$  such that every node outputs “yes”;
- if  $(G, \mathbf{x}) \notin \mathcal{L}$ , then for every certificate  $\mathbf{y}$ , at least one node outputs “no”.

Again, in the (non-anonymous)  $\mathcal{LOCAL}$  model, these two rules must be satisfied for every identity assignment, but the certificates must be the same regardless of the identities. We now recall the class NLD, for *non-deterministic local decision*, as defined in [13]: it is the class of all distributed languages that can be verified in a constant number of rounds in the  $\mathcal{LOCAL}$  model. Similarly, we define  $\text{NLD}^*$ , the anonymous version of NLD, as the class of all distributed languages that can be verified in a constant number of rounds in the anonymous  $\mathcal{LOCAL}$  model. By definition,  $\text{NLD}^* \subseteq \text{NLD}$ .

### 1.3 Our Results

In this paper, we give several evidences supporting the conjecture  $LD^* = LD$ . In particular, we show that it holds for languages defined on paths, with a finite set of input values. More generally, we show that the conjecture holds for *hereditary* languages, that is, languages closed under node deletion. Regarding arbitrary languages, and arbitrary graphs, we prove that the conjecture holds assuming that every node knows an upper bound on the total number of nodes in the input graph. (This upper bound can be arbitrary, and may not be the same for all nodes).

Moreover, we prove that equality between non-anonymous decision and anonymous decision holds in the context of *non-deterministic* local decision, where nodes are given certificates (independent of the identities, if they exist), and the decision consists in verifying these certificates. More precisely, we prove that  $NLD^* = NLD$ . This latter result is obtained by characterizing both NLD and  $NLD^*$ .

### 1.4 Related Work

The question of how to locally decide (or verify) languages has received quite a lot of attention recently. Inspired by classical computation complexity theory, it was suggested in [13] that the study of decision problems may lead to new structural insights also in the more complex distributed computing setting. Indeed, following that paper, which focused on the *LOCAL* model, efforts were made to form a fundamental computational complexity theory for distributed decision problems in various other aspects of distributed computing [13,15,16,17].

The classes LD, NLD and BPLD defined in [13] are the distributed analogues of the classes P, NP and BPP, respectively. The paper provides structural results, developing a notion of local reduction and establishing completeness results. One of the main results is the existence of a sharp threshold for randomization, above which randomization does not help (at least for hereditary languages). More precisely, the BPLD classes were classified into two: below and above the randomization threshold. In [14], the authors show that the hereditary assumption can be lifted if we restrict our attention to languages on path topologies. These two results from [13,14] are used in the current paper in a rather surprising manner. The authors in [14] then “zoom” into the spectrum of classes below the randomization threshold, and defines a hierarchy of an infinite set of BPLD classes, each of which is separated from the class above it in the hierarchy.

The precise knowledge of the number of nodes  $n$  was shown in [13] to be of large impact on non-deterministic decision. Indeed, with such a knowledge every language can be decided non-deterministically in the model of NLD. We note, however, that the knowledge of an arbitrary upper bound on  $n$  (as assumed here in one of our results) seems to be a much weaker assumption, and, in particular, will not suffice for non-deterministically deciding all languages. In the context

of construction problems, it was shown in [28] that in many case, the knowledge of  $n$  (or an upper bound on  $n$ ) is not essential.

The original theoretical basis for non-determinism in local computation was laid by the theory of *proof-labeling schemes* (PLS) [19,24,25,26] originally defined in [26]. As mentioned, this notion resembles the notion of NLD, but differs in the role identities play. Specifically, in PLS the designer of the algorithm may base the certificates' (called labels in the terminology of PLS) construction on the given identity assignment. In contrast, in the model of NLD, the certificates must be the same regardless of the identities of nodes. Indeed, this difference is significant: while every language can be verified by a proof labeling scheme, not every language belongs to NLD [13]. These notions also bear some similarities to the notions of *local computation with advice* [7,10,11,12], *local detection* [1], *local checking* [4], or *silent stabilization* [8]. In addition, as shown later on, the notion of NLD is related also to the theory of *lifts* or *covers* [2,3].

Finally, the classification of decision problems in distributed computing has been studied in several other models. For example, [6] and [23] study specific decision problems in the *CONGEST* model. In [25], the authors study MST verification in the PLS sense but under the *CONGEST* model of communication. In addition, decision problems have been studied in the asynchrony discipline too, specifically in the framework of *wait-free computation* [16,17] and *mobile agents computing* [15]. In the wait-free model, the main issues are not spatial constraints but timing constraints (asynchronism and faults). The main focus of [17] is deterministic protocols aiming at studying the power of the “decoder”, i.e., the interpretation of the results. While this paper essentially considers the AND-checker (since a global “yes” corresponds to all processes saying “yes”), [17] deals with other interpretations, including more values (not only “yes” and “no”), with the objective of designing checkers that use the smallest number of values.

## 2 Deterministic Decision

We conjecture that  $LD = LD^*$ . A support to this conjecture is that it holds for a large class of languages, namely for all *hereditary* languages, that is languages closed under node deletion. For instance, **Coloring** and **MIS** are hereditary, as well as all languages corresponding to hereditary graph families, such as planar graphs, interval graphs, forests, chordal graphs, cographs, perfect graphs, etc.

**Theorem 1.**  $LD^* = LD$  for hereditary languages.

To prove the theorem, it is sufficient to show that  $LD \subseteq LD^*$  for hereditary languages. This immediately follows from the statement and proof of Theorem 3.3 in [13]. Indeed, let  $A$  be a non-anonymous local algorithm deciding  $\mathcal{L}$ . This deterministic algorithm is in particular a randomized algorithm, with success probabilities  $p = 1$  for legal instances, and  $q = 1$  for illegal instance. That is, algorithm  $A$  is a  $(1, 1)$ -decider for  $\mathcal{L}$ , according to the definition in [13]. Since  $\mathcal{L}$  is hereditary, and since  $p^2 + q > 1$ , the existence of  $A$  implies the existence of a

specific deterministic anonymous local algorithm  $D$  for  $\mathcal{L}$ . Indeed, the algorithm  $D$  described in the proof of Theorem 3.3 in [13] is in fact anonymous: it simply collects the ball  $B(u, t)$  of radius  $t$  around each node  $u$  for some constant  $t$ , and  $u$  then decides “yes” or “no” according to whether  $B(u, t) \in \mathcal{L}$  or not, regardless of the identities.

A similar proof, based on Theorem 4.1 in [14], enables to establish the following:

**Theorem 2.**  $\text{LD}^* = \text{LD}$  for languages defined on the set of paths, with a finite set of input values.

Another evidence supporting the conjecture  $\text{LD} = \text{LD}^*$  is that it holds assuming that nodes have access to a seemingly weak oracle. Specifically, this oracle, denoted  $\mathbf{N}$ , simply provides each node with an arbitrarily large upper bound on the total number of nodes in the actual instance. (It is not assumed that all the upper bounds provided to nodes are the same). We denote by  $\text{LD}^{*\mathbf{N}}$  the class of languages that can be decided by an anonymous local algorithm having access to oracle  $\mathbf{N}$ , and we prove the following:

**Theorem 3.**  $\text{LD}^* \subseteq \text{LD} \subseteq \text{LD}^{*\mathbf{N}}$ .

*Proof.* We just need to prove that  $\text{LD} \subseteq \text{LD}^{*\mathbf{N}}$ . Let  $\mathcal{L} \in \text{LD}$ , and let  $A$  be a local (non-anonymous) algorithm deciding  $\mathcal{L}$ . Assume that the running time of  $A$  is  $t$ . We transform  $A$  into an anonymous algorithm  $A'$  deciding  $\mathcal{L}$  in time  $t$ , assuming each node  $u$  in a given input  $G$  has an access to the oracle  $\mathbf{N}$ , i.e., it knows an arbitrary upper bound  $n_u$  on the number of nodes in  $G$ . Algorithm  $A'$  works as follows. Each node  $u$  collects the ball  $B(u, t)$  of radius  $t$  around it. Then, for every possible assignment of identities to the nodes of  $B(u, t)$  taken from the range  $[1, n_u]$ , node  $u$  simulates the behavior of the non-anonymous algorithm  $A$  on the ball  $B(u, t)$  with the corresponding identities. If, in one of these simulations, algorithm  $A$  decides “no”, then  $A'$  decides “no”. Otherwise,  $A'$  decides “yes”.

We now prove the correctness of  $A'$ . If the input  $(G, \mathbf{x}) \in \mathcal{L}$ , then  $A$  accepts it for every identity assignment to the nodes of  $G$ . Therefore, since, for every node  $u$ , every possible identity assignment to the nodes of the ball  $B(u, t)$  can be extended to an identity assignment to all the nodes of  $G$ , all the simulations of  $A$  by  $u$  return “yes”, and hence  $A'$  accepts  $\mathcal{L}$  as well. On the other hand, if  $(G, \mathbf{x}) \notin \mathcal{L}$  then  $A$  rejects it for every identity assignment to the nodes of  $G$ . That is, for every identity assignment to the nodes of  $G$ , at least one node  $u$  decides “no”. (Note that, this node  $u$  may be different for two different identity assignments). Let us fix one identity assignment  $\text{Id}$  to the nodes of  $G$ , in the range  $[1, n]$ , and let  $u$  be one node that decides “no” for  $\text{Id}$ . Let  $B_{\text{Id}}(u, t)$  be the ball  $B(u, t)$  with the identities of the nodes given by  $\text{Id}$ . In  $A'$ , since  $u$  tries all possible identity assignments of the ball  $B(u, t)$  in the range  $[1, n_u]$  with  $n \leq n_u$ , in one of its simulations of  $A$ , node  $u$  will simulate  $A$  on  $B_{\text{Id}}(u, t)$ . In this simulation, node  $u$  decides “no”, and hence algorithm  $A'$  rejects  $\mathcal{L}$  as well.  $\square$



Note that the inclusion  $\text{LD} \subseteq \text{LD}^* \mathbf{N}$  holds when one imposes no restrictions on the individual sequential running time. However, the transformation of a (non-anonymous) local algorithm into an anonymous local algorithm as described in the proof of Theorem 3 is very expensive in terms of individual computation. Indeed, the number of simulations of the original local algorithm  $A$  by each node  $u$  can be as large as  $\binom{n_u}{n_B}$  where  $n_u$  is the upper bound on  $n$  given by the oracle  $\mathbf{N}$ , and  $n_B$  is the number of nodes in the ball  $B(u, t)$ . This bound can be exponential in  $n$  even if the oracle provides a good approximation of  $n$  (even if it gives precisely  $n$ ). It would be nice to establish  $\text{LD} \subseteq \text{LD}^* \mathbf{N}$  by using a transformation not involving a huge increase in the individual sequential computation time.

### 3 Non-deterministic Decision

In the previous section, we have seen several evidences supporting the conjecture that  $\text{LD}^* = \text{LD}$ , but whether it holds or not remains to be proved. In this section, we turn our attention to the non-deterministic variants of these two classes, and show that they coincide. More formally, we have:

**Theorem 4.**  $\text{NLD}^* = \text{NLD}$ .

*Proof.* To prove  $\text{NLD}^* = \text{NLD}$ , it is sufficient to prove  $\text{NLD} \subseteq \text{NLD}^*$ . To establish this inclusion, we provide a sufficient condition for  $\text{NLD}^*$ -membership, and prove that it is a necessary condition for  $\text{NLD}$ -membership.

Let  $I = (G, \mathbf{x})$  and  $I' = (G', \mathbf{x}')$  be two input instances. A *homomorphism* from  $I$  to  $I'$  is a function  $f : V(G) \rightarrow V(G')$  that preserves the edges of  $G$  as well as the inputs to the nodes. Specifically,

$$\{u, v\} \in E(G) \Rightarrow \{f(u), f(v)\} \in E(G'),$$

and  $f$  maps every node  $u \in V(G)$  to a node  $f(u) \in V(G')$  satisfying

$$\mathbf{x}'(f(u)) = \mathbf{x}(u).$$

For instance, assuming the nodes have no inputs, and labeling the nodes of the  $n$ -node cycle  $C_n$  by consecutive integers from 0 to  $n - 1$ , modulo  $n$ , then the map  $f : C_8 \rightarrow C_4$  defined by  $f(u) = u \bmod 4$  is a homomorphism. The trivial map  $g : C_8 \rightarrow K_2$  defined by  $g(u) = u \bmod 2$ , where  $K_2$  is the 2-node clique, is also a homomorphism. To establish conditions for  $\text{NLD}$ - and  $\text{NLD}^*$ -membership, we require the involved homomorphisms to preserve the local neighborhood of a node, and define the notion of *t-local isomorphism*.

Let  $t$  be a positive integer. We say that  $I$  is *t-local isomorphic* to  $I'$  if and only if there exists an homomorphism  $f$  from  $I$  to  $I'$  such that, for every node  $v \in V(G)$ ,  $f$  restricted to  $B_G(v, t)$  is an isomorphism from  $B_G(v, t)$  to  $B_{G'}(f(v), t)$ . We call such a homomorphism  $f$  a *t-local isomorphism*.

Note that a homomorphism is not necessarily a 1-local isomorphism. For instance, the aforementioned map  $f : C_8 \rightarrow C_4$  defined by  $f(u) = u \bmod 4$  is a

1-local isomorphism, but the map  $g : C_8 \rightarrow K_2$  defined by  $g(u) = u \bmod 2$  is not a 1-local isomorphism. To be a 1-local isomorphism, a homomorphism should also insure isomorphism between the balls of radius 1. Also observe that any  $t$ -local isomorphism  $f : G \rightarrow G'$  is onto (because if a node of  $G'$  has no pre-image, then neither do its neighbors have a pre-image, since homomorphisms preserve edges, and so forth). To avoid confusion, it is thus useful to keep in mind that, informally, a  $t$ -local isomorphism goes from a “larger” graph to a “smaller” graph.

**Definition 1.** For positive integer  $t$ , we say that  $\mathcal{L}$  is  $t$ -closed under lift if, for every two instances  $I, I'$  such that  $I$  is  $t$ -local isomorphic to  $I'$ , we have:

$$I' \in \mathcal{L} \Rightarrow I \in \mathcal{L}.$$

So, informally, Definition 1 states that, for a language  $\mathcal{L}$  to be  $t$ -closed under lift, if a “smaller” instance  $I'$  is in  $\mathcal{L}$  then any “larger” instance  $I$  that is a lift of  $I'$ , i.e., satisfying that  $I$  is  $t$ -local isomorphic to  $I'$ , must also be in  $\mathcal{L}$ . The following lemma gives a sufficient condition for  $\text{NLD}^*$ -membership.

**Lemma 1.** Let  $\mathcal{L}$  be a language. If there exists  $t \geq 1$  such that  $\mathcal{L}$  is  $t$ -closed under lift, then  $\mathcal{L} \in \text{NLD}^*$ .

*Proof.* Let  $\mathcal{L}$  be a language, and assume that there exists  $t \geq 1$  such that  $\mathcal{L}$  is  $t$ -closed under lift. We describe an anonymous non-deterministic local algorithm  $A$  deciding  $\mathcal{L}$ , and performing in  $t$  rounds. The certificate of each node  $v$  is a triple  $\mathbf{y}(v) = (i, G', \mathbf{x}')$  where  $G'$  is an  $n$ -node graph with nodes labeled by distinct integers in  $[1, n] = \{1, \dots, n\}$ ,  $i \in [1, n]$ , and  $\mathbf{x}'$  is an  $n$ -dimensional vector. Informally, the certificates are interpreted by  $A$  as follows. The graph  $G'$  is supposed to be a “map” of  $G$ , that is,  $G'$  is interpreted as an isomorphic copy of  $G$ . The integer  $i$  is the label of the node in  $G'$  corresponding to node  $v$  in  $G$ . Finally,  $\mathbf{x}'$  is interpreted as the input of the nodes in  $G'$ .

The algorithm  $A$  performs as follows. Every node  $v$  gets  $B_G(v, t)$ , the ball of radius  $t$  around it; hence, in particular, it collects all the certificates of all the nodes at distance at most  $t$  from it. Then, by comparing its own certificate with the ones of its neighbors, it checks that the graph  $G'$ , and the input  $\mathbf{x}'$  in its certificate, are identical to the ones in the certificates of its neighbors. It also verifies consistency between the labels and the nodes in its ball of radius  $t$ . That is, it checks whether the labels and inputs in the certificate of the nodes in  $B_G(v, t)$  are as described by its certificate. Whenever a node fails to pass any of these tests, it outputs “no”. Otherwise it output “yes” or “no” according to whether  $(G', \mathbf{x}') \in \mathcal{L}$  or not, respectively. (This is doable because we are considering languages that are decidable in the usual sense of sequential computation).

We show that  $A$  performs correctly. If  $(G, \mathbf{x}) \in \mathcal{L}$ , then by labeling the nodes in  $G$  by distinct integers from 1 to  $|V(G)|$ , and by providing the node  $v$  labeled  $i$  with  $\mathbf{y}(v) = (i, G, \mathbf{x})$ , the algorithm  $A$  output “yes” at all nodes, as desired. Consider now a instance  $I = (G, \mathbf{x}) \notin \mathcal{L}$ . Assume, for the purpose of contradiction that there exists a certificate  $\mathbf{y}$  leading all nodes to output “yes”.

Let  $f : V(G) \rightarrow V(G')$  be defined by  $f(v) = i$  where  $i$  is the label of  $v$  in its certificate. Since  $\mathbf{y}$  passes all tests of  $A$ , it means that (1)  $\mathbf{y}(v) = (i, G', \mathbf{x}')$  where the instance  $I' = (G', \mathbf{x}')$  is the same for all nodes, (2)  $f$  restricted to  $B_G(v, t)$  is an isomorphism from  $B_G(v, t)$  to  $B_{G'}(f(v), t)$ , for every node  $v$ , and (3)  $(G', \mathbf{x}') \in \mathcal{L}$ . In view of (2),  $I$  is  $t$ -local isomorphic to  $I'$ . Therefore, (3) implies that  $I = (G, \mathbf{x}) \in \mathcal{L}$ , because  $\mathcal{L}$  is  $t$ -closed under lift. This is in contradiction with the actual hypothesis  $(G, \mathbf{x}) \notin \mathcal{L}$ . Thus, for each certificate  $\mathbf{y}$ , there must exist at least one node that outputs “no”. As a consequence,  $A$  is a non-deterministic algorithm deciding  $\mathcal{L}$ , and thus  $\mathcal{L} \in \text{NLD}^*$ .  $\diamond$

The following lemma shows that the aforementioned sufficient condition for  $\text{NLD}^*$ -membership is a necessary condition for  $\text{NLD}$ -membership.

**Lemma 2.** *Let  $\mathcal{L}$  be a language. If  $\mathcal{L} \in \text{NLD}$ , then there exists  $t \geq 1$  such that  $\mathcal{L}$  is  $t$ -closed under lift.*

*Proof.* Let  $\mathcal{L}$  be a language in  $\text{NLD}$ , and let  $A$  be a non-deterministic (non-anonymous) local algorithm deciding  $\mathcal{L}$ . Assume, for the purpose of contradiction that, for any integer  $t \geq 1$ ,  $\mathcal{L}$  is *not*  $t$ -closed under lift. That is, for any  $t$ , there exist two input instances  $I, I'$  such that  $I$  is  $t$ -local isomorphic to  $I'$ , with  $I \notin \mathcal{L}$  and  $I' \in \mathcal{L}$ . Assume that  $A$  runs in  $t$  rounds. Without loss of generality, we can assume that  $t \geq 1$ . Let  $I = (G, \mathbf{x}) \notin \mathcal{L}$  and  $I' = (G', \mathbf{x}') \in \mathcal{L}$  satisfying  $I$  is  $t$ -local isomorphic to  $I'$ . Since  $I' \in \mathcal{L}$ , there exists a certificate  $\mathbf{y}'$  such that when  $A$  is running on  $I'$  with certificate  $\mathbf{y}'$ , every node output “yes” for every identity assignment. Since  $I$  is  $t$ -local isomorphic to  $I'$ , there exists an homomorphism  $f : I \rightarrow I'$  such that, for every node  $v \in G$ ,  $f$  restricted to  $B_G(v, t)$  is an isomorphism from  $B_G(v, t)$  to  $B_{G'}(f(v), t)$ . Let  $\mathbf{y}$  be the certificate for  $I$  defined by  $\mathbf{y}(v) = \mathbf{y}'(f(v))$ . Consider the execution of  $A$  running on  $I$  with certificate  $\mathbf{y}$ , and some arbitrary identity assignment  $\text{Id}$ .

Since  $A$  performs in  $t$  rounds, the decision at each node  $v$  is taken according to the inputs, certificates, and identities in the ball  $B_G(v, t)$ , as well as the structure of this ball. By the nature of the homomorphism  $f$ , and by the definition of certificate  $\mathbf{y}$ , the structure, inputs and certificates of the ball  $B_G(v, t)$ , are identical to the corresponding structure, inputs and certificates of the ball  $B_{G'}(f(v), t)$ . Balls may however differ in the identities of their nodes. So, let  $v_0$  be the node in  $G$  deciding “no” for  $(G, \mathbf{x})$  with certificate  $\mathbf{y}$ . There exists such a node since  $I \notin \mathcal{L}$ . Let  $v'_0 = f(v_0)$ , and assign the same identities to the nodes in  $B_{G'}(v'_0, t)$  as their corresponding nodes in  $B_G(v_0, t)$ . Arbitrarily extend this identities to an identity assignment  $\text{Id}'$  to the whole graph  $G'$ . By doing so, the two balls are not only isomorphic, but every node in  $B_G(v_0, t)$  has the same input, certificate and identity as its image in  $B_{G'}(v'_0, t)$ . Therefore, the decision taken by  $A$  at  $v_0 \in G$  under  $\text{Id}$  is the same as its decision at  $v'_0 \in G'$  under  $\text{Id}'$ . This is in contradiction to the fact that  $v_0$  decides “no” while  $v'_0$  decides “yes”.  $\diamond$

Lemmas 1 and 2 together establish the theorem.  $\square$

The proof of Lemma 1 also provides an upper bound on the size of the certificates for *graph languages* in  $\text{NLD}$ , that is, for languages in  $\text{NLD}$  with no

input. (This includes, e.g., recognition of interval graphs, and recognition of chordal graphs). Indeed, given  $\mathcal{L} \in \text{NLD}$ , Algorithm  $A$  in the proof of Lemma 1 verifies  $\mathcal{L}$  using a certificate at each node which is essentially an isomorphic copy of the input instance  $(G, \mathbf{x})$ , with nodes labeled by consecutive integers in  $[1, n]$ . If  $\mathcal{L}$  is a graph language, then there is no input  $\mathbf{x}$ , and thus the size of the certificates depends only on the size of the graph. More precisely, we have:

**Corollary 1.** *Let  $\mathcal{L} \in \text{NLD}$  be a graph language. There exists an algorithm verifying  $\mathcal{L}$  using certificates of size  $O(n^2)$  bits at each node of every  $n$ -node graph in  $\mathcal{L}$ .*

We now argue that the above bound is tight, that is, we prove the following.

**Proposition 1.** *There exists a graph language  $\mathcal{L} \in \text{NLD}$  such that every algorithm verifying  $\mathcal{L}$  requires certificates of size  $\Omega(n^2)$  bits.*

*Proof.* Recall that [26] showed that there exists a graph language for which every proof labeling scheme (PLS) requires labels of size  $\Omega(n^2)$  bits (the proof of this latter result appears in a detailed version [27]). Still in the context of PLS, [19] showed that this lower bound holds for two *natural* graph families: specifically, [19] showed that verifying symmetric graphs requires labels of size  $\Omega(n^2)$  bits, and verifying non-3 colorable graphs requires almost the same size of labels, specifically,  $\Omega(n^2 / \log n)$  bits. Note that the certificate size required for verifying a language in NLD is at least as large as the minimum label size required for verifying the language via a proof labeling scheme. Unfortunately, however, one cannot obtain our claim directly from the aforementioned results since it turns out that neither of the two graph languages (namely, symmetric graphs and non-3 colorable graphs) belongs to NLD.

We therefore employ an indirect approach. Specifically, consider a graph  $G$ . We say that  $H$  is a *seed* of  $G$  if there exists a 1-local isomorphism from  $G$  to  $H$ . Suppose  $\mathcal{F}$  is a family of graphs. Let  $\text{Seed-}\mathcal{F}$  denote the family of graphs  $G$ , for which there exists a seed of  $G$  that belongs to  $\mathcal{F}$ . Then, by definition,  $\text{Seed-}\mathcal{F}$  is 1-closed under lift. Indeed, assume that there is a 1-local isomorphism  $g$  from  $G'$  to  $G$ , and let  $H \in \mathcal{F}$  be a seed of  $G$  that belongs to  $\mathcal{F}$ . Then let  $f$  be the 1-local isomorphism from  $G$  to  $H$ . We have that  $f \circ g$  is a 1-local isomorphism from  $G'$  to  $H$ , because, for every  $u \in V(G')$ ,  $B_{G'}(u, 1)$  is isomorphic to  $B_G(g(u), 1)$ , which in turn is isomorphic to  $B_H(f(g(u)), 1)$ . Thus  $H$  is also a seed of  $G'$ .  $\text{Seed-}\mathcal{F}$  is therefore in NLD. Now, in the proof of corollary 2.2 in [27], the authors construct, for every integer  $n$ , a family  $\mathcal{F}_n$  of  $n$ -node graphs that requires proof labels of size  $\Omega(n^2)$ . Note that for every prime integer  $n'$ , a graph  $G$  of size  $n'$  belongs to  $\mathcal{F}_{n'}$  if and only if it belongs to  $\text{Seed-}\mathcal{F}_{n'}$ . Therefore, there exists a graph language, namely,  $\text{Seed-}\mathcal{F}_n$ , that requires certificates of size  $\Omega(n^2)$  bits (at least for prime  $n$ 's).  $\square$

## 4 Conclusion

Again, in this paper, we provide some evidences supporting the conjecture  $\text{LD}^* = \text{LD}$ . For instance, Theorem 3 shows that if every node knows any upper bound

on the number of nodes  $n$ , then all languages in LD can be decided in the anonymous  $\mathcal{LOCAL}$  model. One interesting remark about the  $\mathcal{LOCAL}$  model is that it is guaranteed that at least one node has an upper bound on  $n$ . This is for instance the case of the node with the largest identity. In the anonymous  $\mathcal{LOCAL}$  model, however, there is no such guarantee. Finding a language whose decision would be based on the fact that one node has an upper bound on  $n$  would disprove the conjecture  $\text{LD}^* = \text{LD}$ . Nevertheless, it is not clear whether such a problem exists.

In this paper, we also prove that  $\text{NLD}^* = \text{NLD}$ , that is, our conjecture holds for the non-deterministic setting. It is worth noticing that [13] proved that there exists an NLD-complete problem under the local one-to-many reduction. It is not clear whether such a problem exists for  $\text{NLD}^*$ . Indeed, the reduction in the completeness proof of [13] relies on the aforementioned guarantee that, in the  $\mathcal{LOCAL}$  model, at least one node has an upper bound on  $n$ .

## References

1. Y. Afek, S. Kutten, and M. Yung. The local detection paradigm and its applications to self stabilization. *Theoretical Computer Science*, 186(1-2):199–230, 1997.
2. D. Angluin. Local and Global Properties in Networks of Processors. In *Proc. Twelfth ACM Symp. on Theory of Computing (STOC)*, 82–93, 1980.
3. A. Amit, N. Linial, J. Matousek, and E. Rozenman. Random lifts of graphs. In *Proc. 12th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 883–894, 2001.
4. B. Awerbuch, B. Patt-Shamir, and G. Varghese. Self-Stabilization By Local Checking and Correction. In *Proc. IEEE Symp. on the Foundations of Computer Science (FOCS)*, 1991, 268–277.
5. L. Barenboim and M. Elkin. Distributed  $(\Delta + 1)$ -coloring in linear (in delta) time. In *Proc. 41st ACM Symp. on Theory of computing (STOC)*, 111–120, 2009.
6. A. Das Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg and R. Wattenhofer. Distributed Verification and Hardness of Distributed Approximation. In *Proc. 43rd ACM Symp. on Theory of Computing (STOC)*, 2011.
7. D. Dereniowski and A. Pelc. Drawing maps with advice. *Journal of Parallel and Distributed Computing*, 72:132–143, 2012.
8. S. Dolev, M. Gouda, and M. Schneider. Requirements for silent stabilization. *Acta Informatica*, 36(6), 447–462, 1999.
9. R.G. Gallager, P.A. Humblet, P.M. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Trans. on Programming Languages and Systems*, 5 (1983) 66–77.
10. P. Fraigniaud, C. Gavoille, D. Ilcinkas and A. Pelc. Distributed Computing with Advice: Information Sensitivity of Graph Coloring. In *Proc. 34th Colloq. on Automata, Languages and Programming (ICALP)*, 231–242, 2007.
11. P. Fraigniaud, D. Ilcinkas, and A. Pelc. Communication algorithms with advice. *J. Comput. Syst. Sci.*, 76(3-4):222–232, 2008.
12. P. Fraigniaud, A. Korman, and E. Lebhar. Local MST computation with short advice. In *Proc. 19th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, 154–160, 2007.

13. P. Fraigniaud, A. Korman, and D. Peleg. Local Distributed Decision. In *Proc. 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 708–717, 2011.
14. P. Fraigniaud, A. Korman, M. Parter and D. Peleg. Randomized Distributed Decision. <http://arxiv.org/abs/1207.0252>
15. P. Fraigniaud and A. Pelc. Decidability Classes for Mobile Agents Computing. In *Proc. LATIN 2012: Theoretical Informatics - 10th Latin American Symposium*, 2012.
16. P. Fraigniaud, S. Rajsbaum, and C. Travers. Locality and Checkability in Wait-free Computing. In *Proc. 25th International Symposium on Distributed Computing (DISC)*, 2011.
17. P. Fraigniaud, S. Rajsbaum, and C. Travers. Universal Distributed Checkers and Orientation-Detection Tasks. Submitted, 2012.
18. E. Fusco and A. Pelc. Communication complexity of consensus in anonymous message passing systems. In *Proc. 15th International Conference on Principles of Distributed Systems (OPODIS 2011)*, LNCS 7109, 191–206.
19. M. Göös and J. Suomela. Locally checkable proofs. *Proc. 30th ACM Symp. on Principles of Distributed Computing (PODC)*, 2011.
20. M. Göös, J. Hirvonen, and J. Suomela. Lower bounds for local approximation. In *Proc. 31st Symposium on Principles of Distributed Computing (PODC)*, 2012.
21. H. Hasemann, J. Hirvonen, J. Rybicki, and J. Suomela. Deterministic Local Algorithms, Unique Identifiers, and Fractional Graph Colouring. In *Proc. 19th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, 2012.
22. M. Hanckowiak, M. Karonski, and A. Panconesi. On the Distributed Complexity of Computing Maximal Matchings. *SIAM J. Discrete Math.* 15(1): 41–57 (2001).
23. L. Kor, A. Korman and D. Peleg. Tight Bounds For Distributed MST Verification. In *Proc. 28th Int. Symp. on Theoretical Aspects of Computer Science (STACS)*, 2011.
24. A. Korman and S. Kutten. Distributed verification of minimum spanning trees. *Distributed Computing*, 20:253–266, 2007.
25. A. Korman, S. Kutten, and T. Masuzawa. Fast and Compact Self-Stabilizing Verification, Computation, and Fault Detection of an MST. In *Proc. 30th ACM Symp. on Principles of Distributed Computing (PODC)*, 2011.
26. A. Korman, S. Kutten, and D. Peleg. Proof labeling schemes. *Distributed Computing*, 22:215–233, 2010.
27. A. Korman, S. Kutten, and D. Peleg. Proof labeling schemes. Detailed version. See: <http://ie.technion.ac.il/~kutten/ps-links/ProofLabelingSchemes.ps>
28. A. Korman, J.S. Sereni, and L. Viennot. Toward More Localized Local Algorithms: Removing Assumptions Concerning Global Knowledge. In *Proc. 30th ACM Symp. on Principles of Distributed Computing (PODC)*, 49–58, 2011.
29. F. Kuhn. Weak graph colorings: distributed algorithms and applications. In *Proc. 21st ACM Symp. on Parallel Algorithms and Architectures (SPAA)*, 138–144, 2009.
30. N. Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
31. Z. Lotker, B. Patt-Shamir and A. Rosen. Distributed Approximate Matching. *SIAM J. Comput.* 39(2): 445–460, (2009).
32. M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15:1036–1053, 1986.
33. M. Naor and L. Stockmeyer. What can be computed locally? *SIAM J. Comput.* 24(6): 1259–1277 (1995).

- 34. A. Panconesi and A. Srinivasan. On the Complexity of Distributed Network Decomposition. *J. Algorithms* 20(2): 356–374, (1996).
- 35. D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, 2000.
- 36. D. Seinsche. On a property of the class of  $n$ -colorable graphs. *J. Combinatorial Theory*, Ser. B, 16, pages 191–193, 1974.